

Modeling Brain Circuitry over a Wide Range of Scales

P. Fua and G. Knott*
 EPFL
 CH-1015 Lausanne

Abstract

If we are ever to unravel the mysteries of brain function at its most fundamental level, we will need a precise understanding of how its component neurons connect to each other. Electron Microscopes (EM) can now provide the nanometer resolution that is needed to image synapses, and therefore connections, while Light Microscopes (LM) see at the micrometer resolution required to model the 3D structure of the dendritic network. Since both the topology and the connection strength are integral parts of the brain’s wiring diagram, being able to combine these two modalities is critically important.

In fact, these microscopes now routinely produce high-resolution imagery in such large quantities that the bottleneck becomes automated processing and interpretation, which is needed for such data to be exploited to its full potential. In this paper, we briefly review the Computer Vision techniques we have developed at EPFL to address this need. They include delineating dendritic arbors from LM imagery, segmenting organelles from EM, and combining the two into a consistent representation.

1 Introduction

As our ability to image neurons with light and electron microscopes improves, so does our understanding of their form and function. Today we can image large volumes of both live and fixed brain tissue across a wide range of resolutions. At the micrometer scale, light microscopy (LM) of fluorescently labeled structures reveals dendrites and axons of a subset of neurons that can potentially be reconstructed revealing their complex 3D network, as shown in Fig. 1(b). However, their internal structures and all their surrounding elements remain invisible when using this technique. To see them, one must turn to electron microscopes (EM). These provide images at the nanometer scale making it possible to visualize all the structural elements and especially those that are important for understanding the basic connectivity and activity of different cells. These include synapses, dendritic spines, vesicles, and mitochondria, as depicted by Fig. 1(c).

These recent technologies will therefore provide crucial information about the structural, functional, and plasticity principles that govern neural circuits. And since most neurological and psychiatric disorders involve deviations from these principles, such an understanding is key to treating them. Furthermore, neural circuits exhibit a computational power that no known technology can match. A more

*This work was supported in part by the EU ERC project Micro Nano

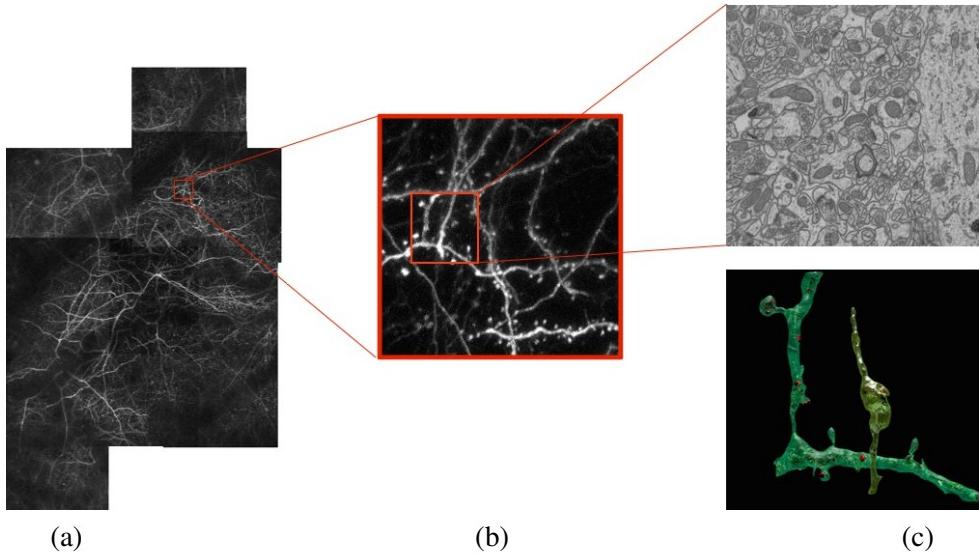


Figure 1: Correlative Microscopy. (a) Fluorescent neurons *in vivo* in the adult mouse brain imaged through a cranial window. (b) Image stack at the $1\mu\text{m}$ resolution acquired using a 2-photon microscope. (c) Image slice of a sub-volume at the 5nm resolution above a reconstruction of a neuron, dendrite, and associated organelles.

thorough understanding of their complexities could therefore spur development of new paradigms and bio-inspired devices that would far outperform existing ones.

However, a major bottleneck stands in the way of this promise: These new microscopes can produce terabytes upon terabytes of image data that is so rich and so complex that humans cannot analyze them effectively in their entirety. In this paper, we will briefly present the algorithms we have developed at EPFL to automatically recover the dendritic and axonal trees, segment intra-neuronal structures from EM images, and register the resulting models. For further details, we refer the interested reader to the original publications.

2 Delineation

The automated delineation of curvilinear structures has been investigated since the inception of the field of Computer Vision in the 1960s and 1970s. Nevertheless, despite decades of sustained effort, full automation remains elusive when the image data is as noisy and the structures exhibit as complex a morphology as they do in microscopy data. As a result, practical systems still require extensive manual intervention that is both time-consuming and tedious. For example, in the DIADEM challenge to map nerve cells, the results of all the finalists still required substantial time and effort to proofread and correct [Ascoli *et al.*, 2010, Peng *et al.*, 2011].

Part of the problem comes from the fact that many existing techniques rely mostly on weak local image evidence, and employ greedy heuristics that can easily get trapped in local minima. As a result, they lack robustness to imaging noise and artifacts. Another common issue is that curvilinear networks are usually treated as tree-like structures without any loops. In practice, however, many interesting networks are not trees since they contain cycles. Furthermore, even among those that really are trees,

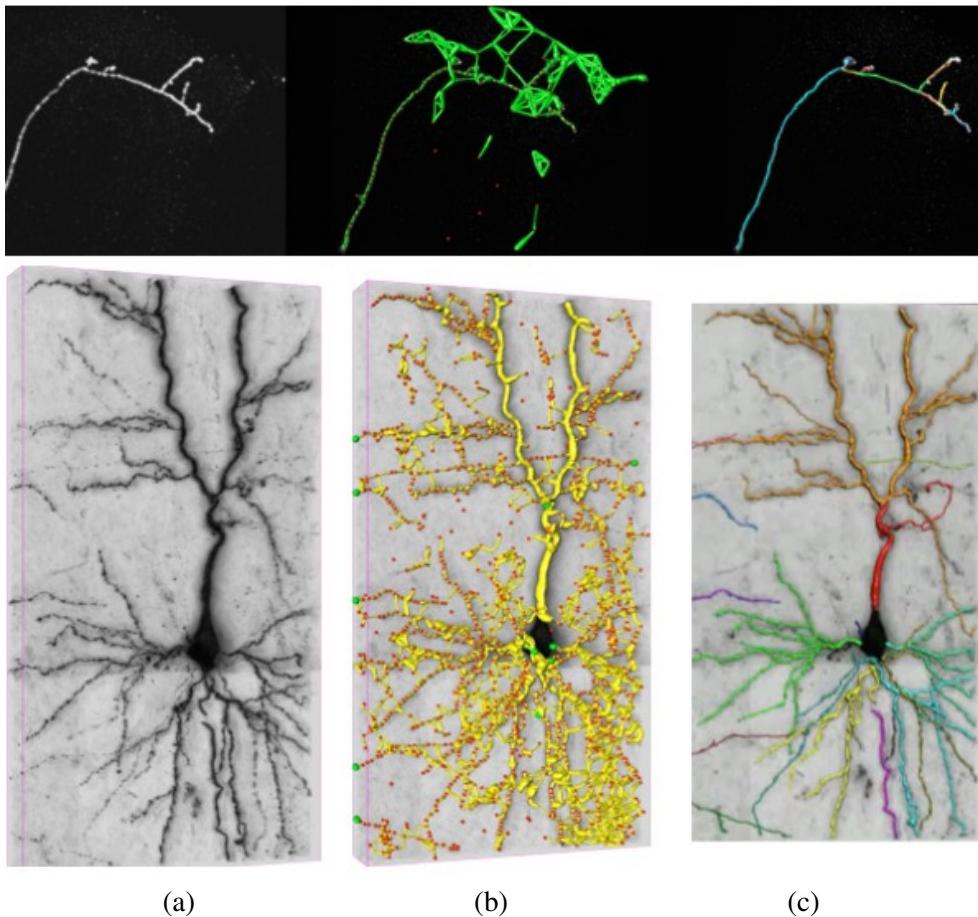


Figure 2: Delineation in confocal (top) and brightfield (bottom) imagery. (a) The original 3D stacks. (b) The nodes appear as red circles with the tubular paths connecting them overlaid in green and yellow. (c) The final 3D delineations.

such as neurites, the imaging resolution is often so low that the branches appear to cross, thus introducing several spurious cycles that can only be recognized once the whole structure has been recovered. In fact, this is reported as one of the major sources of error in [Wang *et al.*, 2011, Chothani *et al.*, 2011, Bas and Erdogmus, 2011, Zhao *et al.*, 2011, Turetken *et al.*, 2011, Choromanska *et al.*, 2012] and a number of heuristics have been proposed to avoid spurious connections in [Chothani *et al.*, 2011, Zhao *et al.*, 2011, Turetken *et al.*, 2011].

2.1 Approach

In our work, we attempt to overcome these limitations by formulating the reconstruction problem as one of solving an Integer Program (IP) on a graph of potential tubular paths. As shown in Fig. 2, the resulting algorithm goes through the following steps:

- We first compute a *tubularity* value at each image location and radius value. It quantifies the likelihood that there exists a tubular structure of this radius at that location. Given an 3D stack, this creates an 4D scale-space tubularity volume.

- We select regularly spaced high-tubularity points as seed points and connect pairs of them that are within a given distance from each other. This results in a directed tubular graph, such as those shown in Fig. 2(b), which serves as an overcomplete representation for the underlying curvilinear networks.
- Having trained a path classifier using such graphs and ground-truth delineations, we assign probabilistic weights to pairs of consecutive edges of a given graph at detection time.
- We use these weights and solve an integer program to compute the maximum-likelihood directed subgraph of this graph to produce a final result such as the one of Fig. 2(c).

These four steps come in roughly the same sequence as those used in most algorithms that build trees from seed points, as in [Fischler *et al.*, 1981, Zhao *et al.*, 2011, Wang *et al.*, 2011, Turetken *et al.*, 2011], but with three key differences. First, whereas heuristic optimization algorithms such as MST followed by pruning or the k-MST algorithm of [Turetken *et al.*, 2011] offer no guarantee of optimality, our approach guarantees that the solution is within a small tolerance of the global optimum. Second, our approach to scoring individual paths using a classifier instead of integrating pixel values as usually done gives us more robustness to image noise and provides peaky probability distributions, which helps ensure that the global optimum is close to the ground truth. Finally, instead of constraining the subgraph to be a tree as many state-of-the-art approaches, we allow it to contain cycles and instead penalize spurious junctions and early branch terminations as described in more details in [Turetken *et al.*, 2012, Turetken *et al.*, 2013a].

2.2 Results

Here, we demonstrate the effectiveness of our approach on the two datasets depicted in Fig. 2:

- *Confocal-Axons*, 8 image stacks of Olfactory Projection Fibers (OPF) of the Drosophila fly acquired using a 3D confocal microscope and taken from the DIADEM competition.
- *Brightfield*: 6 image stacks were acquired by brightfield microscopy from biocytin-stained rat brains.

In both datasets, the neurites form tree structures without cycles. However, in the latter, disjoint branches appear to cross, introducing false loops, due to the low z-resolution. In both cases, we used half the stacks for training and half for testing. We used a semi-automated delineation tool [Turetken *et al.*, 2013b] to extract ground truth tracings from the training stacks and train our path-classifiers.

In Table 1, we compare our approach (OURS) to several state-of-the-art algorithms on the confocal-axons. They are the pruning-based approach (APP2) of [Xiao *et al.*, 2013], the active contour algorithm (OSnake) of [Wang *et al.*, 2011], the NeuronStudio (NS) software of [Wearne *et al.*, 2005], the focus-based depth estimation method (Focus) of [Narayanaswamy *et al.*, 2011], and finally the k-MST technique of [Turetken *et al.*, 2011], the last two of which were finalists in the DIADEM competition. For all these algorithms, we used the implementations provided by their respective authors with default parameters. We report DIADEM scores as described in [Ascoli *et al.*, 2010], which were designed to compare topological accuracy of a reconstructed tree against a ground truth tree.

	OURS	k-MST [Turetken <i>et al.</i> , 2011]	NS [Wearne <i>et al.</i> , 2005]	OSnake [Wang <i>et al.</i> , 2011]	APP2 [Xiao <i>et al.</i> , 2013]
OPF4	0.91	0.87	0.58	0.00	0.67
OPF6	0.91	0.90	0.65	0.80	0.82
OPF7	0.94	0.91	0.42	0.68	0.76
OPF8	0.90	0.74	0.58	0.69	0.63

Table 1: DIADEM scores on four test stacks from the Confocal-Axons dataset. Each row corresponds to an image stack denoted by OP*i*. Higher scores are better.

	BRF1	BRF2	BRF3
OURS	0.05 0.29 0.71 0.65	0.11 0.29 0.81 0.78	0.07 0.28 0.77 0.70
k-MST [Turetken <i>et al.</i> , 2011]	0.10 0.44 0.79 0.88	0.11 0.53 0.84 0.91	0.13 0.35 0.81 0.92
Focus [Narayanaswamy <i>et al.</i> , 2011]	0.39 0.54 0.75 1.00	0.49 0.53 0.90 1.00	0.38 0.46 0.74 1.00
OSnake [Wang <i>et al.</i> , 2011]	0.66 0.63 0.98 0.99	0.66 0.59 0.99 1.00	0.69 0.38 0.95 0.99
APP2 [Xiao <i>et al.</i> , 2013]	0.68 0.64 1.00 1.00	0.63 0.54 1.00 1.00	0.65 0.49 1.00 1.00

Table 2: NetMets scores on the Brightfield dataset. The NetMets software outputs four values for each trial, which are geometric False Positive Rate (FPR), geometric False Negative Rate (FNR), connectivity FPR, and connectivity FNR, respectively from left to right. Lower scores are better.

We also evaluated the APP2 [Xiao *et al.*, 2013], OSnake [Wang *et al.*, 2011], and Focus [Narayanaswamy *et al.*, 2011] algorithms on the Brightfield dataset. Since they do not allow the user to provide multiple root vertices, the DIADEM score of their output cannot be computed. To compare their algorithms to ours, we therefore used the NetMets measure of [Mayerich *et al.*, 2012] instead because it does not rely heavily on roots. As the DIADEM metric, this measure takes as input the reconstruction and the corresponding ground truth tracings. However, it is more local because it does not account for network topology.

Table 2 shows the NetMets scores on the test images of the *Brightfield* dataset. Note that the Focus algorithm of [Narayanaswamy *et al.*, 2011] is specifically designed for brightfield image stacks distorted by a point spread function. Our approach nevertheless brings about a systematic improvement except in one case (BRF3 - connectivity FPR). However, the algorithm does best in this category does significantly worse in the other three.

3 Segmentation

To observe the connectivity between neurons electron microscopy is required. In our work, we have used Focus Ion Beam Scanning Electron Microscopy (FIBSEM) at a 5nm nearly isotropic sampling. The resulting image stacks reveal the fine neuronal structures, including the synaptic contacts. However, segmenting EM data poses unique challenges in part because the volumes are heavily cluttered with structures that exhibit similar textures and are therefore difficult to distinguish based solely on local image statistics. In this section, we outline our approach to segmenting both synapses and mitochondria. They are described in more details in [Becker *et al.*, 2013, Lucchi *et al.*, 2014].

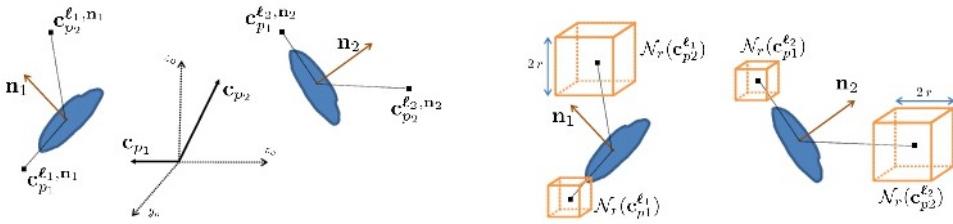


Figure 3: Context features. (a) Relative context cue locations c_p in the global coordinate system x_o, y_o, z_o are rotated according to the orientation estimate of the voxel of interest \mathbf{n}_i to yield locations $c_p^{\ell_i}$ that are consistent. (b) At each of these locations, image channels are summed over cubes of radius r around their center. Our approach employs AdaBoost to select the most discriminative features for synapse segmentation.

3.1 Synapses

3.1.1 Approach

Synapses are difficult to distinguish from other structures based solely on local texture, as shown in Fig. 4. Human experts confirm their presence by looking for nearby post-synaptic densities and vesicles. This protocol cannot be emulated simply by measuring filter responses at the target voxel as in [Kreshuk *et al.*, 2011], pooling features into a global histogram as in [Lucchi *et al.*, 2012, Narasimha *et al.*, 2009] or relying on hand-determined locations for feature extraction as in [Jurrus *et al.*, 2010, Venkataraju *et al.*, 2009].

To emulate this human ability, we designed features we call *context features*, which can be extracted in any cube contained within a large volume centered on the voxel to be classified at 3D location ℓ_i with local orientation \mathbf{n}_i , as depicted in Fig. 3(b). They are computed in several image channels using a number of Gaussian kernels. This yields more than 100,000 potential features and we rely on AdaBoost to select the most discriminative ones.

3.1.2 Results

We evaluated our method on three different EM stacks acquired from different regions of the adult rat brain, the Somatosensory Cortex, the Hippocampus, and the Cerebellum. Example slices from each dataset are shown in Fig. 4 along with our results.

To evaluate the performance of our approach and compare it to that of [Kreshuk *et al.*, 2011], we performed a voxel-wise evaluation against manually acquired ground-truth data. To discount the influence of boundary voxels whose classification may be ambiguous, we defined a testing exclusion zone around the labeled border of the synapse within a distance of d . The voxels within that exclusion zone are ignored and, in Fig. 5, we plot the Jaccard index between the ground-truth labeling and the one the two algorithms produce as a function of d . To highlight the importance of using context, we plot a third curve that corresponds to our approach using only boxes centered on the voxel to be classified, which is much worse than the other two.

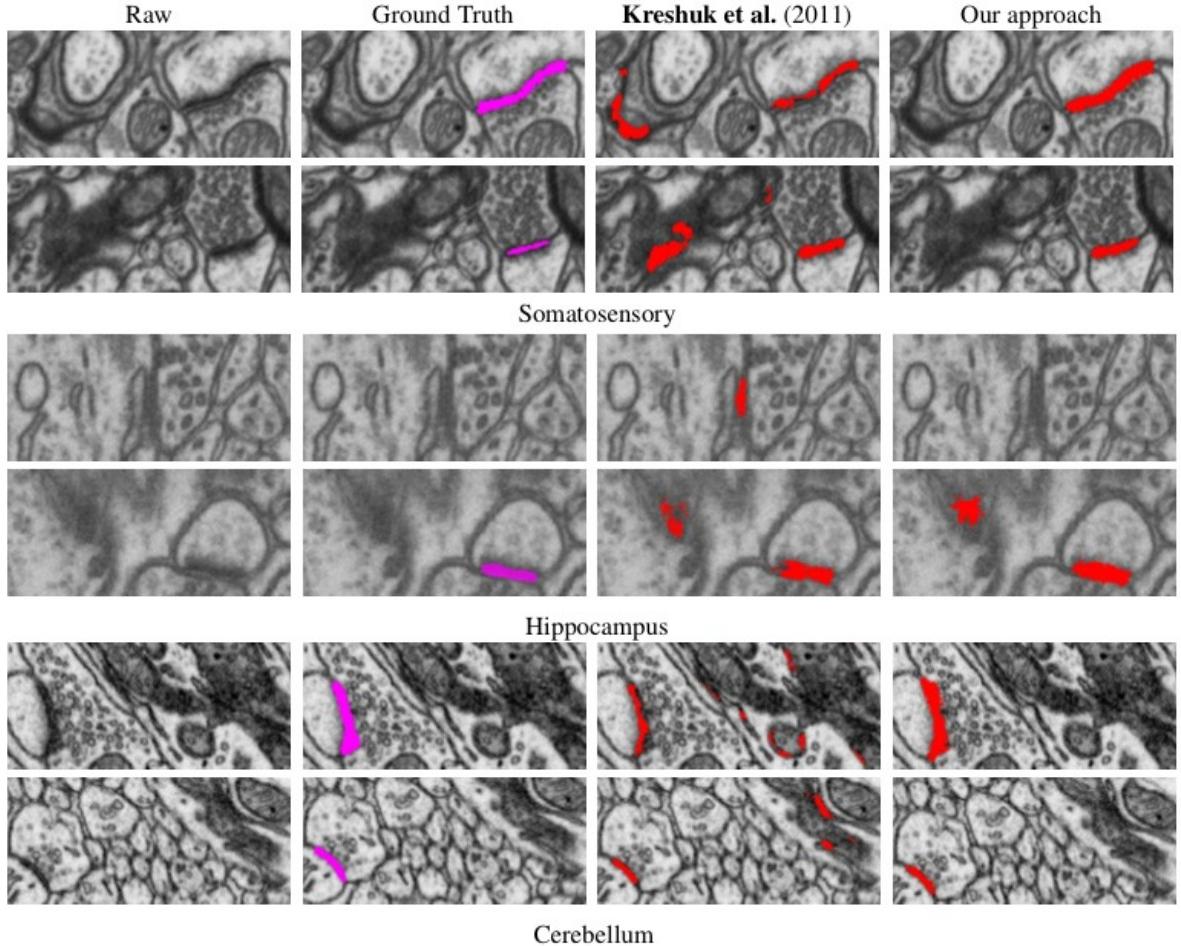


Figure 4: Synapse segmentations overlaid on individual slices from three different datasets after thresholding. Note that our approach yields more accurate results than the competing method with almost no false positives.

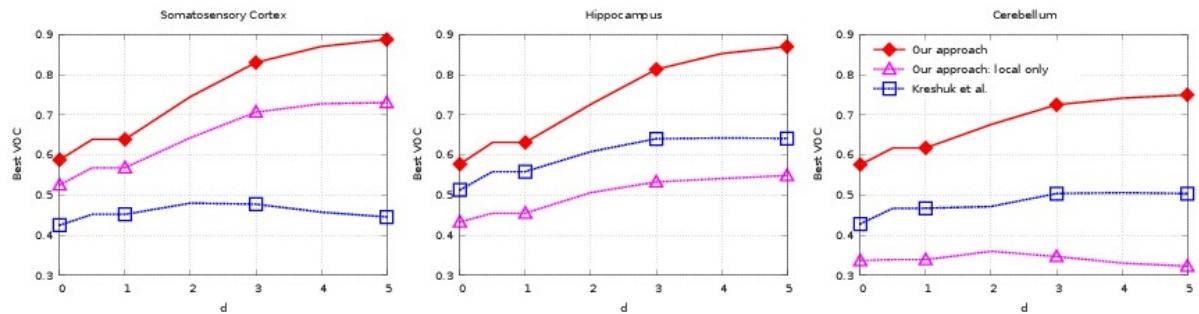


Figure 5: Jaccard index (VOC score) as a function of exclusion zone size d for the different datasets. Our approach outperforms the competing method for all values of d .

3.2 Mitochondria

Mitochondria participate in a wide range of cellular functions and their morphology and localization play a key role in cellular physiology [Campello and Scorrano, 2010]. Furthermore, localization and morphology of mitochondria have been tightly linked to neural functionality. For example, pre- and post-synaptic presence of mitochondria is known to have an important role in synaptic function, as shown in [Lee *et al.*, 2007], and mounting evidence also indicates a close link between mitochondrial function and many neuro-degenerative diseases [Knott *et al.*, 2008, Poole *et al.*, 2008].

New approaches to detecting mitochondria in EM images have therefore begun to appear. For example, in [Vitaladevuni *et al.*, 2008] a Gentle-Boost classifier was trained to detect them based on textural features. In [Narasimha *et al.*, 2009], texton-based mitochondria classification in melanoma cells was performed using a variety of classifiers including k-NN, SVM, and Adaboost. While these techniques achieve reasonable results, they incorporate only textural cues while ignoring shape information. More recently, more sophisticated features have been successfully used in [Sommer *et al.*, 2010, Lucchi *et al.*, 2012, Kumar *et al.*, 2010] in conjunction with either a Random Forest classifier as in [Kreshuk *et al.*, 2011]. The algorithm of [Marquez-Neila *et al.*, 2014] could be used to impose higher-order shape constraints but would be very difficult to extend to 3D volume segmentation because its computational requirements are prohibitive. Our approach overcomes this limitation and extends these earlier techniques by explicitly modeling membranes and exploiting the power of our context features in a Structured SVM framework [Lucchi *et al.*, 2014].

3.2.1 Approach

To reduce the computational complexity, our first step of our approach is to over-segment the image stack into *supervoxels*, that is, small voxel clusters with similar intensities. We use the algorithm of [Achanta *et al.*, 2012] to compute them. It lets us choose their approximate diameter, which we take to be on the order of the known thickness of the outer mitochondrial membranes. This means that membranes are typically one supervoxel thick. All subsequent computations are performed on supervoxels instead of individual voxels, which speeds them up by several orders of magnitude. Our task is now to classify these supervoxels as being inside the mitochondria, part of the membrane, or outside, as shown in Fig. 6(b).

To this end, we introduce a three-class Conditional Random Field (CRF) [Lafferty *et al.*, 2001]. It is defined over a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ whose nodes $i \in \mathcal{V}$ correspond to supervoxels and whose edges $(i, j) \in \mathcal{E}$ connect nodes i and j if they are adjacent in the 3D volume. Each node is associated to a feature vector x_i computed from the image data and a label y_i denoting one of the three classes to which a supervoxel can belong. Let Y be the vector of all y_i , which we will refer to as a *labeling*. The most likely labeling of a volume is then found by minimizing an objective function of the form

$$E^{\mathbf{w}}(Y) = \sum_{i \in \mathcal{V}} D_i^{\mathbf{w}}(y_i) + \sum_{(i,j) \in \mathcal{E}} V_{ij}^{\mathbf{w}}(y_i, y_j), \quad (1)$$

where D_i is referred to as the unary data term and V_{ij} as the pairwise term. The superscript denotes the dependency of these two terms to a parameter vector \mathbf{w} .

The unary data term D_i is taken to be a kernelized function of the context features of Section 3.1.1. The pairwise term is a linear combination of a spatial regularization term and a containment term. The

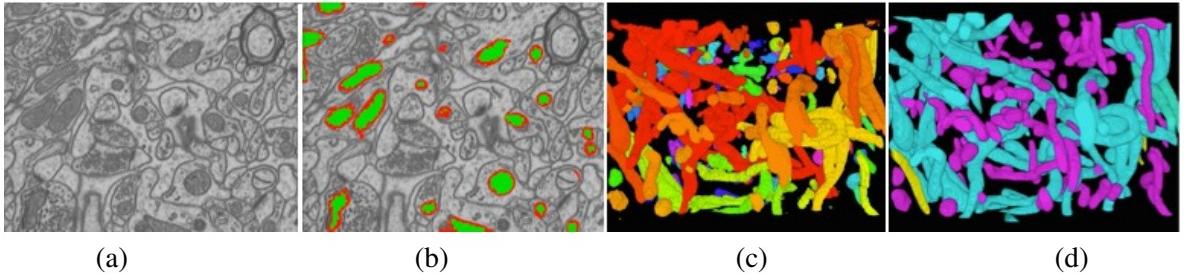


Figure 6: Reconstructed mitochondria. (a) Slice from a 3D image stack. (b) The inside of the mitochondria are overlaid in green and the membranes in red. (c) Raw results. (d) Edited results. The dendritic mitochondria are shown in cyan and axonal ones in purple.

	Seyedhosseini'13	Tsochantaridis'04	Wick'11	Lacoste'13	Ratliff'07	OURS
Hippocampus	83.8%	92.7%	83.3%	92.7%	89.2%	94.8%
Striatum	83.5%	90.6%	89.6%	90.5%	88.1%	92.1%

Table 3: Comparing segmentation performance as measured by the Jaccard index of the foreground class for the Striatum and Hippocampus datasets against that of a number of baselines.

spatial term is learned from data and reflects the transition cost between nodes i and j from label y_i to label y_j . The containment term constrains the membrane class to completely enclose the inside class and to be at least one supervoxel thick, as originally proposed in [Delong and Boykov, 2009]. This containment term is hand-defined and does not depend on any parameters. The set of parameters \mathbf{w} to be learned are therefore the weights given to individual features in the unary term and the spatial regularization term. These parameters are learned within the Structured SVM framework discussed above, which requires solving an inference problem on the supervoxel graph \mathcal{G} .

3.2.2 Results

Fig. 6(c) depicts the 3D reconstructions we obtained from a $3.21 \times m \times 3.21 \mu\text{m} \times 1.08 \mu\text{m}$ volume. In Fig. 6(d), we show the same results after having been proof-read and hand-corrected by a trained neuroscientist. The whole process, including generating the training data, took a little under 2 hours. For comparison purposes, the neuroscientist re-generated these results entirely manually and that took him about 6 hours for a similar level of precision in terms of the mitochondria volumes and surface areas, which are the relevant biological quantities. In other words, automation reduced the required amount of manual intervention by a factor 3. Going further will require deploying tools based on deformable models such as those of [Neuenschwander *et al.*, 1994, Neuenschwander *et al.*, 1997, Jorstad and Fua, 2014] to automatically refine mitochondria boundaries and break apart incorrectly merged ones.

To further quantify the performance of our approach, we compared it against other recent automatic methods on image stacks from the Hippocampus and Striatum, which are similar to those we used to detect synapses. In Table 3, we report the Jaccard index for the foreground and membrane class jointly, which is representative for this task since whole mitochondria are the object of interest being segmented. The first one is a very recent mitochondria segmentation method [Seyedhosseini *et al.*, 2013] that does *not* rely on structured learning. Instead, it trains a cascade of classifiers at different scales and has been shown to outperform earlier algorithms based on Neural Networks, SVMs, and Random Forests on EM

imagery. The others correspond to different approaches to performing structured learning. As can be seen, we consistently outperform the competing methods.

4 Registration

Registering LM and EM stacks such as those of Fig. 1(b,c) is required to identify the same region in both images and to combine the specific information each modality provides, as discussed earlier. However, this is challenging because the scale-discrepancy between the two modalities— 1000nm for EM vs 5nm for LM—produces drastic appearance changes. It makes it impractical to use standard registration techniques that rely on maximizing image similarity, such as those described in [Pluim *et al.*, 2003].

Instead, we have proposed in [Serratell *et al.*, 2014] a new approach for matching graph structures embedded in 3D volumes, which can deal with the scale-change while being robust to topological differences between the two graphs and even changes in the distances between vertices, unlike earlier graph-matching techniques such as those of [Deng *et al.*, 2010, Smeets *et al.*, 2010]. It requires no initial position estimate, can handle non-linear deformations, and does not rely on local appearance or global distance matrices. Instead, given graphs extracted from the two images or image-stacks to be registered, we treat graph nodes as the features to be matched. We model the geometric mapping from one data set to the other as a Gaussian Process whose predictions are progressively refined as more correspondences are added. These predictions are in turn used to explore the set of all possible correspondences starting with the most likely ones, which allows convergence at an acceptable computational cost even though no appearance information is available.

4.1 Approach

Given graphs $\mathcal{G}^A = (\mathbf{X}^A, \mathbf{E}^A)$ and $\mathcal{G}^B = (\mathbf{X}^B, \mathbf{E}^B)$ extracted from image-stacks A and B , let the \mathbf{E} s denote edges and the \mathbf{X} s nodes. The edges, in turn, are represented by dense sets of points forming 3D paths connecting the nodes. Our goal is to use these two graphs to find a geometrical mapping m from A to B such that $m(\mathbf{x}_i^A)$ is as close as possible to \mathbf{x}_j^B in the least-squares sense assuming that \mathbf{x}_i^A and \mathbf{x}_j^B are corresponding voxels.

If correspondences between points belonging to the two graphs were given, we could directly use the Gaussian Process Regression (GPR) as in [Rasmussen and Williams, 2006] to estimate a non-linear mapping that would yield a prediction of m and its associated variance. In our case, however, the correspondences are initially unavailable and cannot be established on the basis of local image information because the A and B are too different in appearance. In short, this means that we must rely only on geometrical properties to simultaneously establish the correspondences and estimate the underlying non-linear transform. Since attempting to do this directly for all edge points would be computationally intractable, our algorithm goes through the following two steps:

- 1. Coarse alignment:** We begin by only matching graph nodes so that the resulting mapping is a combination of an affine deformation and a smooth non-linear deformation. We initialize the search by randomly picking D correspondences, which roughly fixes relative scale and orientation, and using them to instantiate a Gaussian Process (GP). We then recursively refine it as follows: Given some matches between \mathcal{G}^A and \mathcal{G}^B nodes, the GP serves to predict where other \mathcal{G}^A nodes should map and restricts the set of potential correspondences. Among these possibilities, we

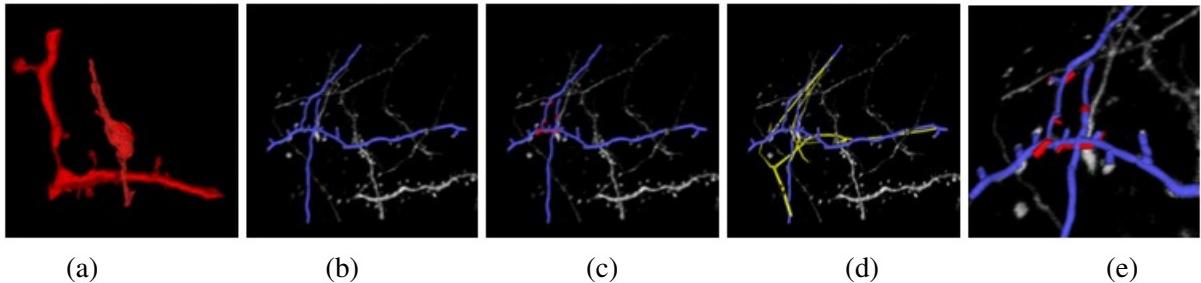


Figure 7: **Light and electron microscopy neuronal trees.** (a) Graph structure extracted from the electron microscopy image stack, in red. (b) Segmented light microscope neurons in blue. (c) After the non-linear registration process using ATS-RGM, the EM segmented neuron is deformed and aligned over the LM extracted neuron. (d) Registration using CPD, in yellow, which falls into a local minimum. (e) A zoom over the region where the EM stack has been extracted. The two neurons have been completely aligned. Best viewed in color.

select the most promising one and use it to refine the GP. Repeating this procedure recursively until enough mutually consistent correspondences have been established and backtracking when necessary lets us quickly explore the set of potential correspondences and recover an approximate geometric mapping.

2. Fine alignment: Having been learned only from potentially distant graph nodes, the above-mapping is coarse. To refine it, we also establish correspondences between points that form the edges connecting the nodes in such a way that distances along these edges, which we will refer to as *geodesic* distances, are changed as little as possible between the two graphs. Because there are many more such points than nodes, this would be extremely expensive to do from scratch. Therefore, we constrain the correspondence candidates to edges between already matched nodes and rely on the Hungarian algorithm of [Munkres, 1957] to perform the optimal assignment quickly.

4.2 Results

Fig. 7 illustrates the two stages of our approach applied to the EM and LM stacks of Fig. 1. Even though the two images look extremely different, our algorithm returns a non-rigid deformation that lets us correctly superpose the two stacks. The technique is generic and allows us to correctly align other biological structures, such as blood-vessels networks, that are non-linearly transformed and extracted with different techniques, without having to pre-aligning them and in a manageable amount of time.

5 Conclusion

If we are ever to unravel the mysteries of brain function at its most fundamental level, we will need a precise understanding of how neurons connect to each other. With the advent of new high-resolution light and electron microscopes, fast computers, and high-capacity storage media, the data required to perform this task is now becoming available. Electron microscopes (EM) can now provide the nanometer resolution that is needed to image synapses, and therefore connections, while Light Microscopes (LM) see at the micrometer resolution required to model the 3D structure of the dendritic network. Since both the arborescence and the connections are integral parts of the wiring diagram, combining these two

modalities is critically important to answer a growing need for automated quantitative assessment of neuron morphology and connectivity.

Here, we have reviewed our approach to addressing this daunting task. Our algorithms are effective at delineating linear structures in LM, segmenting mitochondria and synapses in EM, and putting the results into a unified coordinate systems to produce a joint representation.¹ However, we have so far only modeled small fractions of cells, which only represent minute parts of simple neural circuit. Our challenge therefore is now to scale up our methods so that they can handle much larger volumes, which will involve parallelizing them and using GPUs, instead of CPUs, to massively increase the processing speed.

References

- [Achanta *et al.*, 2012] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Suesstrunk. SLIC Superpixels Compared to State-Of-The-Art Superpixel Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, November 2012.
- [Ascoli *et al.*, 2010] G. A. Ascoli, K. Svoboda, and Y. Liu. Digital Reconstruction of Axonal and Dendritic Morphology Diadem Challenge, 2010. <http://diademchallenge.org/>.
- [Bas and Erdogmus, 2011] E. Bas and D. Erdogmus. Principal Curves as Skeletons of Tubular Objects - Locally Characterizing the Structures of Axons. *Neuroinformatics*, 9(2-3):181–191, 2011.
- [Becker *et al.*, 2013] C. Becker, K. Ali, G. Knott, and P. Fua. Learning Context Cues for Synapse Segmentation. *IEEE Transactions on Medical Imaging*, 32(10):1864–1877, October 2013.
- [Campello and Scorrano, 2010] S. Campello and L. Scorrano. Mitochondrial Shape Changes: Orchestrating Cell Pathophysiology. *EMBO Reports*, 11(9):678–84, 2010.
- [Choromanska *et al.*, 2012] A. Choromanska, S.F. Chang, and R. Yuste. Automatic Reconstruction of Neural Morphologies with Multi-Scale Graph-Based Tracking. *Frontiers in Neural Circuits*, 6(25), 2012.
- [Chothani *et al.*, 2011] P. Chothani, V. Mehta, and A. Stepanyants. Automated Tracing of Neurites from Light Microscopy Stacks of Images. *Neuroinformatics*, 9:263–278, 2011.
- [Delong and Boykov, 2009] A. Delong and Y. Boykov. Globally Optimal Segmentation of Multi-Region Objects. In *International Conference on Computer Vision*, pages 285–292, 2009.
- [Deng *et al.*, 2010] K. Deng, J. Tian, J. Zheng, X. Zhang, X. Dai, and M. Xu. Retinal Fundus Image Registration via Vascular Structure Graph Matching. *International Journal of Biomedical Imaging*, 2010.
- [Fischler *et al.*, 1981] M.A Fischler, J.M. Tenenbaum, and H.C. Wolf. Detection of Roads and Linear Structures in Low-Resolution Aerial Imagery Using a Multisource Knowledge Integration Technique. *Computer Vision, Graphics, and Image Processing*, 15(3):201–223, March 1981.

¹Some of the corresponding software can be downloaded from our lab’s website <http://cvlab.epfl.ch/software>

- [Jorstad and Fua, 2014] A. Jorstad and P. Fua. Refining Mitochondria Segmentation in Electron Microscopy Imagery with Active Surfaces. In *ECCV Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment*, 2014.
- [Jurrus *et al.*, 2010] E. Jurrus, A. Paiva, S. Watanabe, J. Anderson, R. Whitaker, B.W. Jones, R. Marc, and T. Tasdizen. Detection of Neuron Membranes in Electron Microscopy Images Using a Serial Neural Network Architecture. *Medical Image Analysis*, 14(6):770–783, 2010.
- [Knott *et al.*, 2008] A. Knott, G. Perkins, R. Schwarzenbacher, and E. Bossy-Wetzel. Mitochondrial Fragmentation in Neurodegeneration. *Nature Reviews Neuroscience*, 9(7):505–18, 2008.
- [Kreshuk *et al.*, 2011] A. Kreshuk, Christoph N. Straehle, C. Sommer, U. Koethe, G. Knott, and F.A. Hamprecht. Automated Segmentation of Synapses in 3D EM Data. In *International Symposium on Biomedical Imaging*, Chicago, IL, USA, 2011.
- [Kumar *et al.*, 2010] R. Kumar, A. Vazquez-Reina, and H. Pfister. Radon-Like Features and Their Application to Connectomics. In *Workshop on Mathematical Methods in Biomedical Image Analysis*, San Francisco, CA, USA, 2010.
- [Lacoste-Julien *et al.*, 2013] S. Lacoste-Julien, M. Jaggi, M. Schmidt, and P. Pletscher. Block-Coordinate Frank-Wolfe Optimization for Structural SVMs. In *International Conference on Machine Learning*, pages 53–61, Atlanta, GA, USA, 2013.
- [Lafferty *et al.*, 2001] J. Lafferty, A. Mccallum, and F. Pereira. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *International Conference on Machine Learning*, pages 282–289, Bellevue, WA, USA, 2001.
- [Lee *et al.*, 2007] D. Lee, K. Lee, W. Ho, and S. Lee. Target Cell-Specific Involvement of Presynaptic Mitochondria in Post-Tetanic Potentiation at Hippocampal Mossy Fiber Synapses. *The Journal of Neuroscience*, 27(50):13603–13, 2007.
- [Lucchi *et al.*, 2012] A. Lucchi, K. Smith, R. Achanta, G. Knott, and P. Fua. Supervoxel-Based Segmentation of Mitochondria in EM Image Stacks with Learned Shape Features. *IEEE Transactions on Medical Imaging*, 31(2):474–486, Feb 2012.
- [Lucchi *et al.*, 2014] A. Lucchi, Y. Li, C. Becker, P. Marquez-Neila, K. Smith, G. Knott, and P. Fua. Learning Structured Models for Segmentation of 2D and 3D Imagery. *IEEE Transactions on Medical Imaging*, 2014. In Press.
- [Marquez-Neila *et al.*, 2014] P. Marquez-Neila, P. Kohli, C. Rother, and L. Baumela. Non-parametric Higher-order Random Fields for Image Segmentation. In *European Conference on Computer Vision*, pages 269–284, Zurich, Switzerland, 2014.
- [Mayerich *et al.*, 2012] D. Mayerich, C. Bjornsson, J Taylor, and B. Roysam. Netmets: Software for Quantifying and Visualizing Errors in Biological Network Segmentation. *BMC Bioinformatics*, 13, 2012.
- [Munkres, 1957] J. Munkres. Algorithms for the Assignment and Transportation Problems. *Journal of the Society of Industrial and Applied Mathematics*, 5(1):32–38, March 1957.

- [Narasimha *et al.*, 2009] R. Narasimha, H. Ouyang, A. Gray, S.W. McLaughlin, and S. Subramaniam. Automatic Joint Classification and Segmentation of Whole Cell 3D Images. *Pattern Recognition*, 42:1067–1079, 2009.
- [Narayanaswamy *et al.*, 2011] A. Narayanaswamy, Y. Wang, and B. Roysam. 3D Image Pre-Processing Algorithms for Improved Automated Tracing of Neuronal Arbors. *Neuroinformatics*, 9(2-3):219–231, 2011.
- [Neuenschwander *et al.*, 1994] W. Neuenschwander, P. Fua, G. Székely, and O. Kubler. Initializing Snakes. In *Conference on Computer Vision and Pattern Recognition*, pages 613–615, June 1994.
- [Neuenschwander *et al.*, 1997] W. Neuenschwander, P. Fua, G. Székely, and O. Kubler. Velcro Surfaces: Fast Initialization of Deformable Models. *Computer Vision and Image Understanding*, 2(65):237–245, February 1997.
- [Peng *et al.*, 2011] H. Peng, F. Long, T. Zhao, and E. Myers. Proof-Editing is the Bottleneck of 3D Neuron Reconstruction: the Problem and Solutions. *Neuroinformatics*, 9(2):103–105, 2011.
- [Pluim *et al.*, 2003] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever. Mutual Information Based Registration of Medical Images: A Survey. *IEEE Transactions on Medical Imaging*, 22(8):986–1004, 2003.
- [Poole *et al.*, 2008] A. Poole, R. Thomas, L. Andrews, H. McBride, A. Whitworth, and L. Pallanck. The Pink1/parkin Pathway Regulates Mitochondrial Morphology. *Proceedings of the National Academy of Sciences of the United States of America*, 105(5):1638–43, 2008.
- [Rasmussen and Williams, 2006] C. E. Rasmussen and C. K. Williams. *Gaussian Process for Machine Learning*. MIT Press, 2006.
- [Ratliff *et al.*, 2007] N.D. Ratliff, J. Andrew Bagnell, and M.A. Zinkevich. (Online) Subgradient Methods for Structured Prediction. In *International Conference on Artificial Intelligence and Statistics*, pages 380–387, San Juan, Puerto Rico, 2007.
- [Serratell *et al.*, 2014] E. Serratell, M.A. Pinheiro, R. Sznitman, J. Kybic, F. Moreno-Noguer, and P. Fua. Non-Rigid Graph Registration Using Active Testing Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2014.
- [Seyedhosseini *et al.*, 2013] M. Seyedhosseini, M. Sajjadi, and T. Tasdizen. Image Segmentation with Cascaded Hierarchical Models and Logistic Disjunctive Normal Networks. In *International Conference on Computer Vision*, pages 2168–2175, Sydney, Australia, 2013.
- [Smeets *et al.*, 2010] D. Smeets, P. Bruyninckx, J. Keustermans, D. Vandermeulen, and P. Suetens. Robust Matching of 3D Lung Vessel Trees. In *MICCAI Workshop on Pulmonary Image Analysis*, Beijing, China, 2010.
- [Sommer *et al.*, 2010] C. Sommer, C. Straehle, U. Koethe, and F. Hamprecht. Interactive Learning and Segmentation Tool Kit. In *Systems Biology of Human Disease*, pages 230–33, 2010.
- [Tschantaridis *et al.*, 2004] I. Tschantaridis, T. Hofmann, T. Joachims, and Y. Altun. Support Vector Machine Learning for Interdependent and Structured Output Spaces. In *International Conference on Machine Learning*, page 104, 2004.

- [Turetken *et al.*, 2011] E. Turetken, G. Gonzalez, C. Blum, and P. Fua. Automated Reconstruction of Dendritic and Axonal Trees by Global Optimization with Geometric Priors. *Neuroinformatics*, 9(2-3):279–302, 2011.
- [Turetken *et al.*, 2012] E. Turetken, F. Benmansour, and P. Fua. Automated Reconstruction of Tree Structures Using Path Classifiers and Mixed Integer Programming. In *Conference on Computer Vision and Pattern Recognition*, Providence, RI, USA, June 2012.
- [Turetken *et al.*, 2013a] E. Turetken, F. Benmansour, B. Andres, H. Pfister, and P. Fua. Reconstructing Loopy Curvilinear Structures Using Integer Programming. In *Conference on Computer Vision and Pattern Recognition*, Portland, OR, USA, June 2013.
- [Turetken *et al.*, 2013b] E. Turetken, F. Benmansour, and P. Fua. Semi-Automated Reconstruction of Curvilinear Structures in Noisy 2D Images and 3D Image Stacks. Technical report, EPFL, 2013.
- [Venkataraju *et al.*, 2009] K.U. Venkataraju, A Paiva, E. Jurrus, and T. Tasdizen. Automatic Markup of Neural Cell Membranes Using Boosted Decision Stumps. In *IEEE Symposium on Biomedical Imaging: From Nano to Macro*, pages 1039–42, Boston, MA, USA, 2009.
- [Vitaladevuni *et al.*, 2008] S.N. Vitaladevuni, Y. Mishchenko, A. Genkin, D. Chklovskii, and Kristen Harris. Mitochondria Detection in Electron Microscopy Images. In *Workshop on Microscopic Image Analysis with Applications in Biology*, New York, NY, USA, 2008.
- [Wang *et al.*, 2011] Y. Wang, A. Narayanaswamy, and B. Roysam. Novel 4D Open-Curve Active Contour and Curve Completion Approach for Automated Tree Structure Extraction. In *Conference on Computer Vision and Pattern Recognition*, pages 1105–1112, Colorado Springs, CO, USA, 2011.
- [Wearne *et al.*, 2005] S.L. Wearne, A. Rodriguez, D.B. Ehlenberger, A.B. Rocher, S.C. Henderson, and P.R. Hof. New Techniques for Imaging, Digitization and Analysis of Three-Dimensional Neural Morphology on Multiple Scales. *Neuroscience*, 136(3):661–680, 2005.
- [Wick *et al.*, 2011] M. Wick, K. Rohanianesh, K. Bellare, A. Culotta, and A. Mccallum. Samplerank: Training Factor Graphs with Atomic Gradients. In *International Conference on Machine Learning*, pages 777–784, Bellevue, WA, USA, 2011.
- [Xiao *et al.*, 2013] C. Xiao, M. Staring, Y. Wang, D. Shamonin, and B.C. Stoel. Multiscale Bi-Gaussian Filter for Adjacent Curvilinear Structures Detection with Application to Vasculature Images. *IEEE Transactions on Image Processing*, 22(1):174–188, 2013.
- [Zhao *et al.*, 2011] T. Zhao, J. Xie, F. Amat, N. Clack, P. Ahammad, H. Peng, F. Long, and E. Myers. Automated Reconstruction of Neuronal Morphology Based on Local Geometrical and Global Structural Models. *Neuroinformatics*, 9:247–261, May 2011.